

20 May 2014

Dr. Harold Hawkins
ONR Code 341
Office of Naval Research
875 North Randolph SL
Arlington, VA 22203-1995

Reference: US Navy Contract N00014-12-C-0653: "The Model Analyst's Toolkit: Scientific Model Development, Analysis, and Validation"
Charles River Analytics Contract No. C12186

Subject: Contractor's Quarterly Status Report #7
Reporting Period: 20-February-2014 to 19-May-2014

Dear Dr. Hawkins,

Please find enclosed 1 copy of the Quarterly Status Report for the referenced contract. Please feel free to contact me with any questions regarding this report or the status of the "The Model Analyst's Toolkit: Scientific Model Development, Analysis, and Validation" effort.

Sincerely,



W. Scott Neal Reilly
Principal Investigator

cc: Cheryl Gonzales, DCMA
Annetta Burger, ONR
Whitney McCoy, Charles River Analytics

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 20 MAY 2014		2. REPORT TYPE		3. DATES COVERED 00-00-2014 to 00-00-2014	
4. TITLE AND SUBTITLE The Model Analyst's Toolkit:Scientific Model Development, Analysis, and Validation				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Charles River Analytics Inc.,625 Mount Auburn St.,Cambridge,MA,02138				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 16	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Charles River Analytics

Monthly Technical Progress Report No. R12186-07

Reporting Period: February 20, 2014 to May 19, 2014

Government Contract No. N00014-12-C-0653

Charles River Analytics Contract No. C12186

The Model Analyst's Toolkit: Scientific Model Development, Analysis, and Validation Quarterly Status Report

Principal Investigator: Scott Neal Reilly

Charles River Analytics
625 Mount Auburn Street
Cambridge, MA 02138
617-491-3474

May 20, 2014

Distribution Statement A: Approved for public release; distribution is unlimited.

The views, opinions, and findings contained in this report are those of the authors and should not be construed as an official Agency position, policy, or decision, unless so designated by other official documentation.

1. Executive Summary

The proposed research effort builds on and extends the work of the previous ONR-funded “Validation Coverage Toolkit for HSCB Models” project. The overall objectives of the ongoing research program are:

- Help scientists create, analyze, refine, and validate rich scientific models
- Help computational scientists verify the correctness of their implementations of those models
- Help users of scientific models, including decision makers within the US Navy, to use those models correctly and with confidence
- Use a combination of human-driven data visualization and analysis, automated data analysis, and machine learning to leverage human expertise in model building with automated analyses of complex models against large datasets

Specific objectives for the current effort include:

- **Fluid temporal correlation analysis.** Our objective is to design a new method for performing temporally fluid correlation analysis for temporal sets of data and implement the method as a new prototype component within the Model Analyst’s Toolkit (MAT) software application.
- **Automated suggestions for model construction and refinement.** Our objective is to design and implement a prototype mechanism that learns from data how factors interact in non-trivial ways in scientific models.
- **Data validation and repair.** Our objective is to design and implement a prototype capability to identify likely errors in data based on anomalies relative to historic data and to use models of historic data to offer suggested repairs.
- **System prototyping.** Our objective is to incorporate all improvements into the MAT software application and make the resulting application available to the government and academic research community for use in scientific modeling projects.
- **Evaluation of applicability to multiple scientific domains.** Our objective is to ensure (and demonstrate) that MAT can be applied to a wide range of scientific domains by identifying and building at least one neurological and/or physiological model and analyze the associated data with MAT, making any extensions to the MAT tool that are needed to support the analysis of such a model.

2. Overview of Problem and Technical Approach

2.1. Summary of the Problem

One of the most powerful things scientists can do is to create models that describe the world around us. Models help scientists organize their theories and suggest additional experiments to run. Validated models also help others in more practical applications. For instance, in the hands of military decision makers, human social cultural behavior (HSCB) models can help predict instability and the socio-political effects of missions, whereas models of the human brain and

mind can help educators and trainers create curricula that more effectively improve the knowledge, skills, and abilities of their pupils.

While there are various software tools that are used by the scientific community to help them develop and analyze their models (e.g., Excel, R, Simulink, Matlab), they are largely so general in purpose (e.g., Excel, R) or so focused on computational models in particular (e.g., Simulink, Matlab), that they are not ideal for rapid model exploration or for use by non-computational scientists. They also largely ignore the problem of validating the models, especially when the models are positing causal claims as most interesting scientific models do. To address this gap, Charles River Analytics undertook the “Validation Coverage Toolkit for HSCB Models” project with ONR. Under this effort, we successfully designed, implemented, informally evaluated, and deployed a tool called the Model Analyst’s Toolkit (MAT), which focused on supporting social scientists to visualize and explore data, develop causal models, and validate those models against available data (Neal Reilly, 2010; Neal Reilly, Pfeffer, Barnett et al., 2011, 2010).

As part of the development of the MAT tool, we identified four important extensions to that research program that would further support the scientific modeling process:

- Correlation analyses are still the standard way of identifying relationships between factors in a model, but correlations are fundamentally flawed as a tool for analyzing potentially causal or predictive relationships as they assume instantaneous effects. Even performing correlation analyses with a temporal offsets between streams of data is insufficient as the temporal gap between the causal or predictive event and the following event may not be the same every time (either because of variability in the system being modeled or because of variability introduced by a fixed sampling rate). What we need is a novel way of evaluating the true predictive power across streams of data that can deal with fluid offsets between changes in one stream of data and follow events in the other stream of data.
- Modeling complex phenomena is a fundamentally difficult task. Human intuition and analysis is by far the most effective way of performing this task, but even humans can be overwhelmed by the complexity of modeling the systems they are studying (e.g., socio-political system, human neurophysiology). Automated tools, while not especially good at generating reasonable scientific hypotheses, *are* extremely good at processing large amounts of data. We believe there is an opportunity for computational systems to enhance human scientific inquiry. Under the “Validation Coverage Toolkit for HSCB Models” project, we demonstrated how automated tools could help human scientists to analyze and validate their models against data. We believe a similar approach can be used to help suggest modifications to the human-built models to make them better match the available data. To be useful, however, such automated analyses will need to be rich enough to suggest subtle data interactions that are most likely to be missed by the human scientist. For instance, correlations (especially correlations that take into account fluid temporal displacements) could be used to identify likely relationships between streams of data, but such an approach would miss complex, non-linear relationships between interrelated factors that cannot be effectively analyzed with

simple two-way correlations. For instance, if crime waves are associated with increases in unemployment *or* drops in the police presence, that would be hard to identify with a correlation analysis. We need richer automated data analysis techniques that can extract complex, non-linear, multi-variable relationships between data if we are to effectively suggest model improvements to human scientists.

- Even if a scientific model is sound, if the data sets provided as inputs to the model are unreliable, the results of the model are still suspect. And, unfortunately, data will often be wrong. For instance, HSCB surveys are notoriously unreliable and biased for a variety of reasons, and neurological and physiological data can be corrupted by broken or improperly used sensors. If it were possible to identify when data was unreliable and, ideally, even repair the data, then the models that are using the data could once again be effectively used.
- The MAT tool we developed under the “Validation Coverage Toolkit for HSCB Models” project was focused primarily on assisting social scientists in the analysis, refinement, and validation of HSCB models. In parallel with that effort, however, we also took an opportunity to apply MAT to evaluating neurological and physiological data under the DARPA-funded CRANIUM (Cognitive Readiness Agents for Neural Imaging and Understanding Models) program. We discovered the generality of the MAT tool makes it potentially applicable to a great number of different scientific domains. MAT proved to be a useful, but peripheral tool, in CRANIUM. We believe MAT could be applied to a broader suite of scientific modeling problems than it has been so far.

2.2. Summary of our Approach

To address these identified gaps and opportunities, we are extending MAT’s support for model development, analysis, refinement, and validation; enhancing MAT to analyze and repair data; and demonstrating MATs usefulness in additional scientific modeling domains. Our approach encompasses the following four areas, which correspond to the four gaps/opportunities identified in the previous section:

- **Temporally Fluid Correlation Analysis.** We are designing a new method to perform Temporally Fluid Correlational Analysis on temporal sets of data, and we are implementing the method as a new component within the MAT software application. The version of MAT at the beginning of the new effort supported correlation analysis for temporally offset data; it shifts the two data streams being compared by a fixed offset that is based on the sampling rate of the data (i.e., data that is sampled annually will be shifted by one year at a time), performs a standard correlation on the shifted data, plots the correlation value against the amount of the offset, and then repeats the process for the next offset amount. If two data streams are shifted by a fixed offset (e.g., changes in one stream are always followed by a comparable value in the other stream after a fixed time), then this method will find that offset. Under the current effort, we are expanding on this capability to support fluid temporal shifts within the data streams. That is, we are making it possible to identify when the temporal offset between the

change in the first data stream and its effect in the second stream is not a static amount of time.

- **Automated suggestions for model construction and refinement.** We are designing and implementing a mechanism to learn how factors interact in non-trivial ways in scientific models. In particular, we are developing a method for learning disjuncts, conjuncts, and negations. This mechanism starts with the model developed by the scientist user and make recommendations for possible adjustments to make it more complete by performing statistical data mining and machine learning.
- **Data validation and repair.** Recognizing that data contains errors is plausible once we understand the relationships between data sets. That is, if we are able to develop models of the correlations between sets of data, then we can build systems that notice when these correlations do not hold in new data, indicating possible errors in data. For instance, if we know that public sentiment tends to vary similarly between nearby towns, then when one town shows anomalous behavior, we can reasonably suspect problems with the data. There might be local issues that cause the anomaly, but it is, at least, worth noting and bringing to the attention of the user of the data and model. As MAT is designed to help analyze models and recognize inter-data relationships, it is primed to perform exactly this analysis. Existing methods perform similar types of analysis for environmental data (Dereszynski & Dietterich, 2007, 2011). For instance, a broken thermometer can be identified and the data from it even estimated by looking at the temperature readings of nearby thermometers, which will generally be highly correlated.
- **Application to multiple scientific modeling domains.** To ensure (and demonstrate) that MAT can be applied to a wide range of scientific domains, we are identifying and building at least one neurological and/or physiological model and analyzing the associated data with MAT, making any extensions to the MAT tool that are needed to support the analysis of such a model. The initial MAT effort focused on HSCB models; by focusing this effort on harder-science models at much shorter time durations, we believe we can effectively evaluate an interesting range of applications of the MAT tool.

3. Current Activities and Status

During the current reporting period, we made progress on the causal model recommendation component, the new data synthesis component, and the feature learning component. We have also begun a quality assurance (QA) effort to ensure the constantly developing MAT system continues to be stable for our increasing user based.

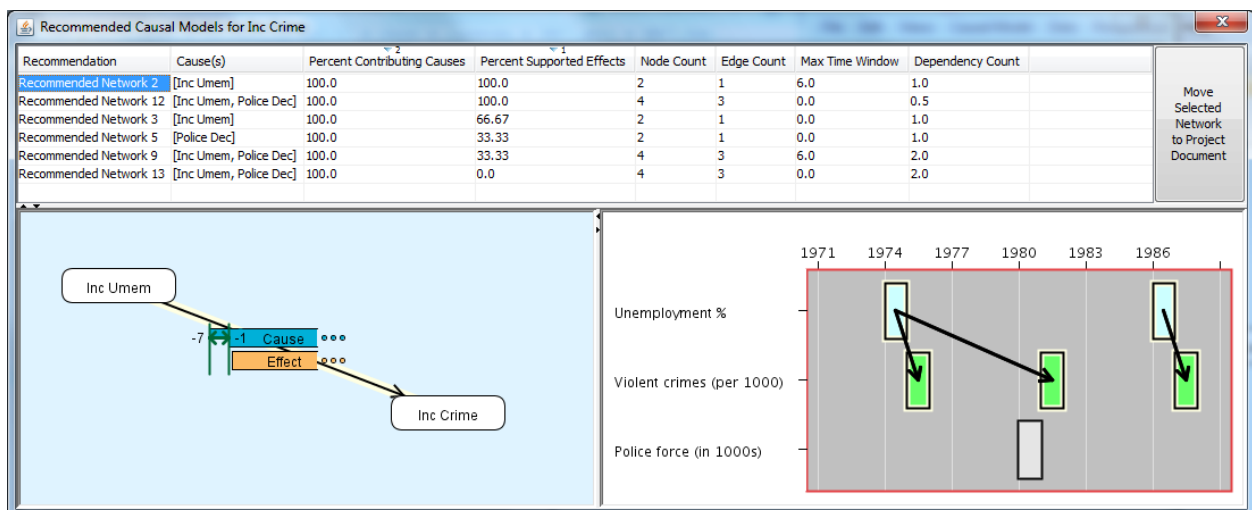
3.1. Causal Model Recommendation Improvements

The causal model recommender automatically suggests modifications to user-defined causal models from the available data. It has been improved to return a set of suggested model edits and display the results in a user friendly way. The recommendations are now a list of causal models where no model is strictly dominated by another model in the list. This eliminates any obviously worse off choices, but also makes no assumptions about the user's preferences in

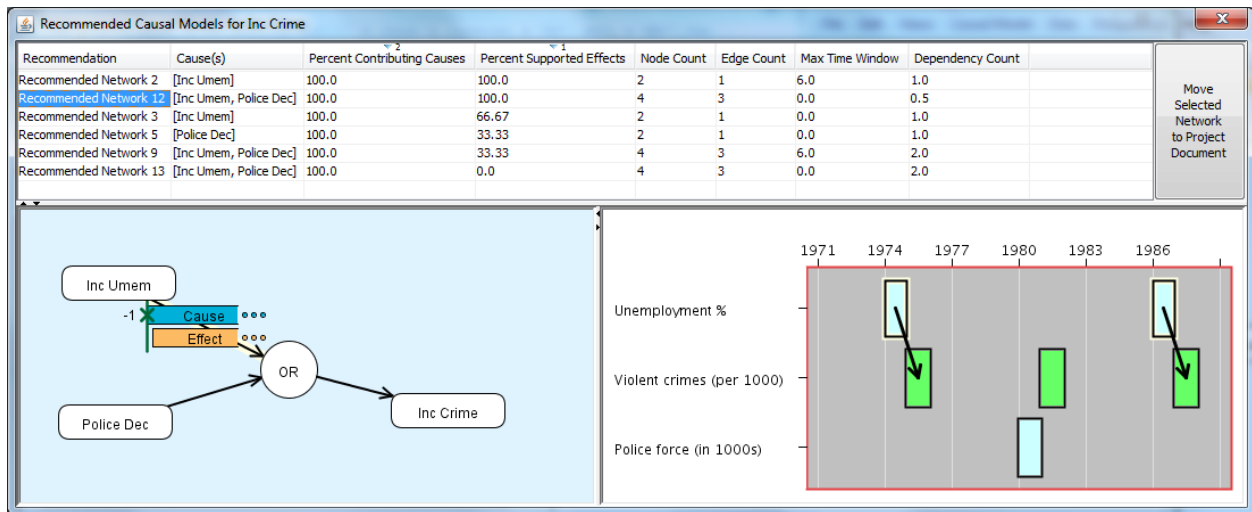
tradeoffs regarding various characteristics of the causal models. Characteristics of the causal models used to determine if a causal model is included in the Pareto Frontier include: performance (number of supported effects and contributing causes), model size (number of nodes and edges, where simpler models are preferred), and temporal aspects (size of temporal window, where models that use temporally closer causes and effects are preferred). The recommendations include simple causal models that only have a single cause for the effect of interest, but more complex causal models are also generated where multiple causes are combined using logic nodes. We have chosen to only make these single-level recommendations as we believe it will result in the most plausible and acceptable modifications to the models created by the scientist-user.

MAT currently uses two algorithms for generating causal model recommendations. The first examines all possible combinations of causes with all possible combinations of temporal offsets for the data provided. This approach quickly becomes computationally expensive, so a second approach is also included where more complex models are built using the results from simpler models and thereby, greatly reducing the number of possible causal models to evaluate. However, this approach may miss a causal model (e.g., models with multiple causes) that is found by the first approach. Also, the recommender displays a progress bar during the operation and the user can cancel it if it is taking too long or if a model is found that seems acceptable or interesting.

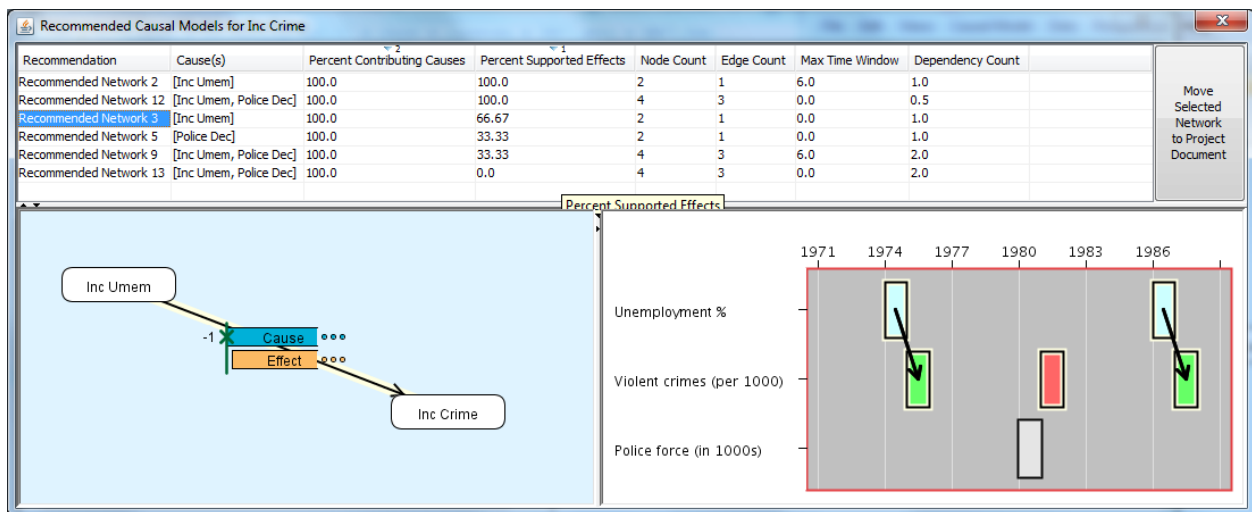
Both of these algorithms return a Pareto Frontier of causal models, which prevents any obviously inferior models from being presented to the user, but there can still be many recommendations generated. Therefore, the recommender results are displayed in a sortable table where each row is a causal model recommendation and the user can sort based on the aspects of the causal models that are most important to them by clicking on the table's column headers. The recommendation table makes it easy for the user to explore the various causal models and see how it influences model validation. For example, in the following screenshot, the user can pick between a simple causal model with a larger temporal window:



Or a more complex causal model with a smaller temporal window:



Causal models that have lower performance, but excel in other aspects are also included in the table. For example, a simple causal model with a small temporal window may be preferred to either of the previous two models even though it does not provide support for all of the effects:



3.2. MAT Data Synthesis Capability

This period, we also made progress on the new data synthesis capability that has been requested by users. This capability is presented to the user as a new tab in MAT and lets the user create new data series by manipulating and combining existing data series. So, for instance, a new data series can be created that is the average (or max, or sum, or...) of other data series, and this new data series can be analyzed for its relationship to other events in the data. This will allow to express (and learn) causal models like, “whenever the sum of the percentage of people unemployed and people who are unhappy with their job crosses a threshold....”

During the current reporting period we completed the design and implementation of a new internal software infrastructure to support this new capability and also improved the design of

the data synthesis capability. One improvement is the modification to fundamental operations. Testing of the functionality revealed that some types of operations, like multiplication and division, were better grouped together. One key problem has been finding a way to indicate the order of non-commutative operations like subtraction in the UI. Since the predecessor nodes in a graph are undifferentiated, there is no obvious way to indicate subtraction in a single step. The current design uses the properties pane to indicate which inputs are subtracted versus added:

Operation:	Multiplication	▼
*	▼	Irrigation Factor
*	▼	Normalized Precipitation
/	▼	Inflation Annualized

New graph properties design showing display for an operator-type node in data synthesis.

In this design, multiplication and division use the same operator which has a “*/÷” label, and likewise addition and subtraction use the same operator. When inputs are connected to the operator node they automatically appear in the properties box as shown above. The user can then select either “*” or “/” from a drop-down list. This design obviated operators such as inversion. In the new design, inversion is achieved by using a multiply/divide operator and then selecting “/” from the drop down.

We also completed the logic for the data synthesis evaluation during the reporting period. Since the synthesis graph can be complex and have many dependencies, the system must determine the order of evaluation of nodes, and must also must validate each operation and verify that there are no cycles. When the user selects a node to synthesize and presses the “Generate Data” button, the system validates the network and generates the resulting synthetic dataset represented by that node or gives an error describing the validation failure, if any.

3.1. Automatic Feature Extraction

In many domains, causal models can often be more readily described as patterns of qualitative features rather than quantitative relationships. In MAT, users can identify qualitative features in data streams that represent meaningful events, such as “spikes in crime.” The existing feature recognition system uses these user-identified events as exemplars in a learning-by-example approach, automatically searching for repeated, temporal patterns of these events in the data.

This only works, however, when the user knows which features are of interest ahead of time. We expect this often be the case, but not always, so we are including functionality in MAT to automatically mine the available data for “interesting” features that have explanatory power

with respect to explaining causes of other (user-defined) events. To provide this capability, we have been developing an automated approach to extracting features in data streams by using a non-linear optimization algorithm, the Nelder-Mead Simplex algorithm, to identify structural, qualitative features of a data series. This algorithm divides a time series into the optimal combination of structural features using the featurization “language” (from Olszewski, 2001) discussed in previous reports (see Figure 1).

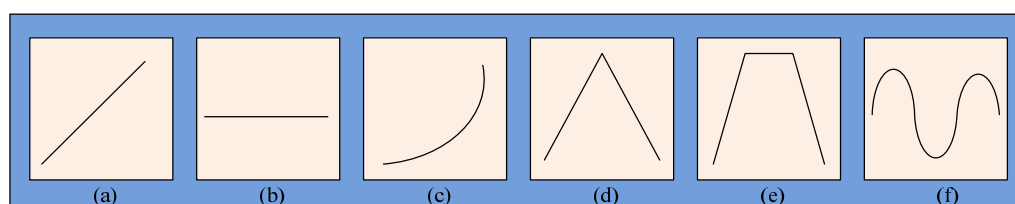
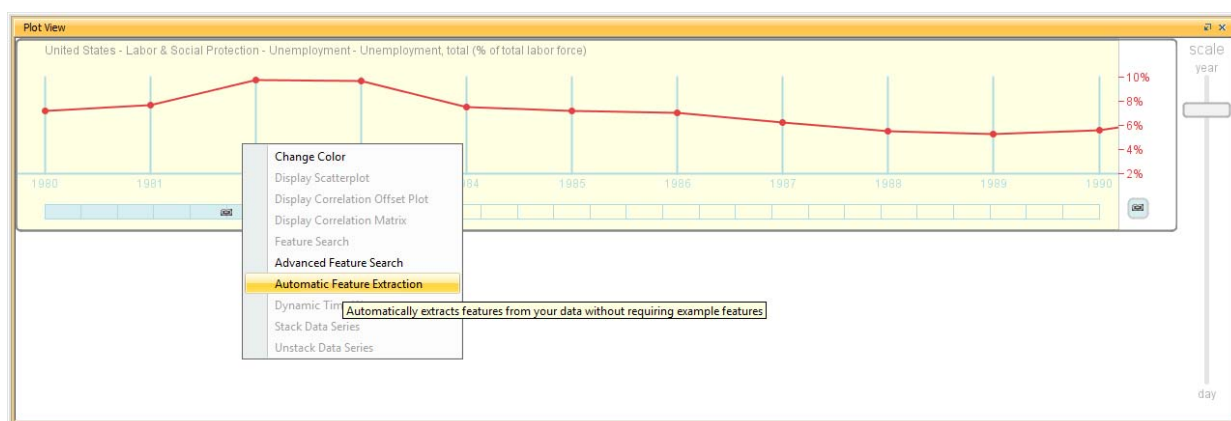


Figure 1. Six common function morphologies that can comprise qualitative features: (a) slope, (b) constant, (c) exponential, (d) triangle, (e) trapezoidal, (f) sinusoidal

When features are selected by the automatic feature extraction algorithm, they are then clustered into *meaningful* concepts. For example, similarly shaped exponential increases in crime are grouped together in a concept called “increases in crime.” Currently, this mechanism is based on the morphologies given in Figure 1, but we plan to explore additional clustering algorithms that can group features at a finer granularity according to the parameters of their structural representation and their duration over time. This new capability has been fully integrated into the MAT user interface (see Figure 2).



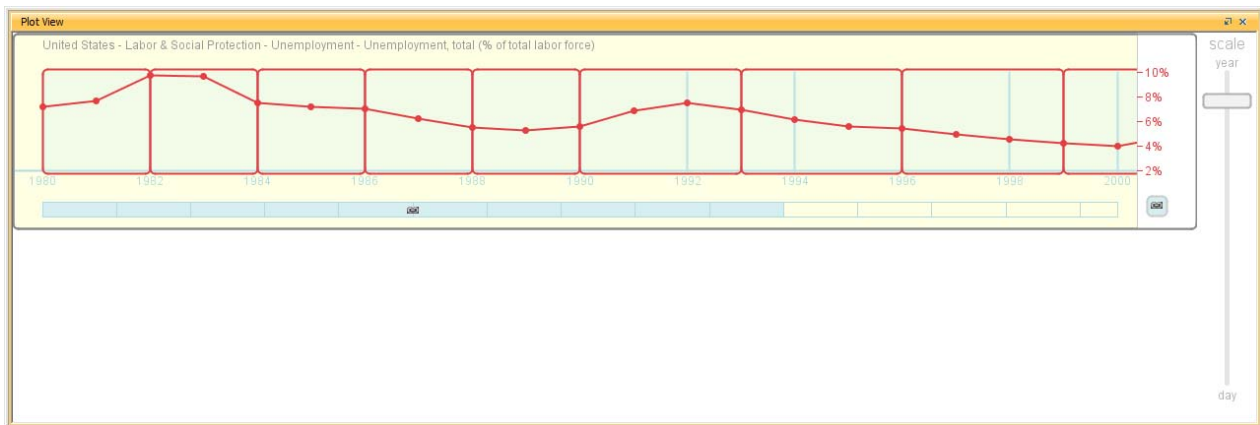


Figure 2 Automatic feature extraction identifies qualitative structures in a time series

This approach will fully featurize a data stream, which will often generate more features than are useful or interesting. Therefore, we are combining this automatic feature extraction with a heuristic version of the TF-IDF (term frequency-inverse document frequency) algorithm from document analysis to identify features that are not only characteristic of a time series (e.g., frequent in the data stream, but infrequent in other data streams), but also those that are uncommon, but quite extreme and meaningful from a causal modeling perspective (e.g., you might only have one stock market crash in your data, so it isn't frequent, but it is still extreme enough to be interesting).

In MAT, this automated feature extraction can be used in conjunction with the causal model recommender, providing additional candidate causes that may not have already been identified by the user. With this capability, MAT now provides the user with novel suggestions of causal relationships based on features that might otherwise have been overlooked, assisting users in refining and validating their causal models.

3.2. MAT Quality Assurance

As we deploy MAT to more users and continue to extend and modify the codebase, we are finding it necessary to devote some effort to ensuring the robustness of the software. To this end, Quality Assurance testing is ongoing as MAT. This includes regression testing of existing features of MAT as well as testing of new features and improvements as they are integrated into the new version of MAT. As part of this testing process a User Acceptance test is being compiled for use as a regression script for future releases.

4. Planned Activities

During the upcoming reporting period, we plan to focus on the following tasks:

- Implementing the heuristic addition to our TF-IDF feature analysis functionality that will pull out features that represent uncommon but large events in the data.
- Completing the implementation of the data synthesis capability.

- Preparing for and presenting at the annual ONR program review.
- Presenting MAT to Dr. Adam Russell of IARPA and exploring possible uses of MAT on projects of interest to IARPA.

5. Evaluation and Transition

We continue to focus on making MAT available to the government and academic research communities and to look for opportunities to use MAT on a variety of ongoing research efforts.

To support this effort, during the current reporting period we worked with Erin Fitzgerald to include a write-up on MAT in a MINERVA program email. This resulted in follow-up discussions and software deliveries to Dr. Dominick' Wright and Joint Advanced Warfighting Division (JAWD) and Dr. Adam Russell at IARPA. Dr. Russell will be visiting Charles River on June 9 and we will give him a demo of MAT and discuss his interest in MAT to support his ongoing efforts.

We also found out this period that our ADAPTER SBIR program with AFRL/RH has been selected to go to Phase II. MAT is being used on ADAPTER to analyze neuro-physiological data from cyber operators to evaluate cognitive workload during team-based cyber operations.

Table 1 summarizes our progress in this regard to date. We will continue to update this table as we make additional progress and will include it as a regular part of future status reports.

Program	Customer	Comments
On-going efforts		
Tourniquet Master Trainer (TMT) (Phase I SBIR)	US Army's Telemedicine & Advanced Technology Research Center (TATRC)	MAT is being used to visualize and analyze data from sensors on a medical manikin that indicate whether a number of novel medical devices used to combat junctional and inguinal hemorrhaging are being applied properly. This program is about to begin a Phase II where MAT will continue to be used both by Charles River Analytics and our partners at the University of Wisconsin.

Laparoscopic Surgery Training System (LASTS) (Phase II SBIR)	US Navy's Office of Naval Research (ONR)	Under lasts, Charles River and Caroline Cao at Wright State University are using MAT to analyze data collected from the location of the laproscopic surgery tools tools during an experiment. Surgical tools are instrumented with markers and 3D data is collected on their location as the person performs the task. This is an ongoing Phase II SBIR program.
Cognitive Readiness Agents for Neural Imaging and Understanding Models (CRANIUM) (Phase I SBIR)	US Navy's Office of Naval Research (ONR)	MAT was used to visualize and extract patterns of stress and workload from neuro-physiological data for training systems. This was a Phase I SBIR program that did not progress to Phase II.
Business Intelligence Visualization for Organizational Understanding, Analysis, and Collaboration (BIVOUAC) Phase II SBIR	US Navy's Space and Naval Warfare Systems Command (SPAWAR)	MAT is being evaluated as part of the BIVOUAC SBIR program, which provides data analysis and visualization for Enterprise Resource Planning (ERP) systems for the Navy. This is an ongoing Phase II SBIR program.
Adaptive toolkit for the Assessment and augmentation of Performance by Teams in Real time (ADAPTER) (Phase I SBIR)	US Air Force Research Lab Human Effectiveness Directorate (AFRL/RH)	MAT is being used to analyze neuro-physiological data from cyber operators to evaluate cognitive workload during team-based cyber operations. This program has been chosen to go to Phase II and we awaiting contract award.
Anticipated Efforts		

Enhancing Intuitive Decision Making Through Implicit Learning (I2BRC) (ONR Basic Research Challenge BAA)	US Navy's Office of Naval Research (ONR) Charles River is a subcontractor to DSCI MESH Solutions, LLC	The intention is to use MAT to help analyze neuro-physiological data to help better understand how implicit learning and intuitive decision making work. This is an ongoing BAA program, though no data has yet been collected to analyze.
A system for augmenting training by Monitoring, Extracting, and Decoding Indicators of Cognitive Load (MEDIC)	US Army's Telemedicine & Advanced Technology Research Center (TATRC)	We are evaluating the practicability of using MAT to analyze and visualize neuro-physiological data from combat medic trainees to identify periods of stress and cognitive overload. This is a SBIR Phase I program where MAT is being evaluated. The Phase II proposal is currently being written.
Soldier's Intelligence Fusion Toolkit (SIFT)	US Army Research Laboratory (ARL)	Extend MAT for ARL research objective in high-level information fusion, exploitation, social network analysis and knowledge management research. A BAA white paper submission has been requested and has been submitted.

Table 1. MAT Transition and Use Progress

In addition we have provided copies of MAT to the following institutions based on their requests for the software: the University of Michigan, Arizona State University, Kansas State University, University of California at Los Angeles, the Naval Medical Research Unit at Wright Patterson Air Force Base, Concordia University (Montreal), the University of Wisconsin, and the Air Force Research Laboratory's Human Effectiveness Directorate, the Intelligence Advanced Research Projects Agency (IARPA), and the Joint Advanced Warfighting Division (JAWD).

Finally, during the previous reporting period, we submitted a paper abstract on using MAT for data-driven model refinement and validation to the American Political Science Association that has been approved for a presentation at the annual conference in August.

6. Budget and Project Tracking

As of April 30, 2014, we have spent \$554,396, or 60% of our total budget of \$928,224, in 55% of the scheduled time. Our current funding is \$662,477, so we have spent 84% of our available funding.

We anticipate spending a bit quickly over the next month to support the IARPA demo and the ONR annual program review and then scaling back our effort a bit to ensure we stay on track to ensure the current funding increment lasts through September 20 per instructions from ONR's contracts office.

Overall, we believe we are in good shape to complete the project on time and on budget.

7. References

- Dereszynski, E. W. & Dietterich, T. G. (2011). Spatiotemporal models for data-anomaly detection in dynamic environmental monitoring campaigns. *ACM Transactions on Sensor Networks (TOSN)*, 8.
- Dereszynski, E. & Dietterich, T. (2007). Probabilistic models for anomaly detection in remote sensor data streams.
- Granger, C. W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, 424-438.
- Granger, C. W. (1980). Testing for causality: a personal viewpoint. *Journal of Economic Dynamics and Control*, 2, 329-352.
- Manning, C., Raghavan, R., and Schütze, H. (2008). *Introduction to Information Retrieval*.: Cambridge University Press.
- Myers, C. & Rabiner, L. (1981). A level building dynamic time warping algorithm for connected word recognition. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 29, 284-297.
- Neal Reilly, S. (2010). *Validation Coverage Toolkit for HSCB Models*. Charles River Analytics Inc.
- Neal Reilly, W. S., Pfeffer, A., and Barnett, J. (2010). A Metamodel Description Language for HSCB Modeling. In *Advances in Cross-Cultural Decision Making*.
- Neal Reilly, W. S., Pfeffer, A., Barnett, J., Chamberlain, J., and Casstevens, R. (2011). A Computational Toolset for Socio-Cultural Data Exploration, Model Refinement, and Model Validation. In *Human Social Culture Behavior (HSCB) Focus*.
- Olszewski, R. (2001). *Generalized feature extraction for structural pattern recognition in time-series data*. Carnegie Mellon University, School of Computer Science.

Salvador, S.& Chan, P. (2007). Toward accurate dynamic time warping in linear time and space. *Intelligent Data Analysis*, 11, 561-580.

Sugihara, G., May, R., Ye, H., Hsieh, C., Deyle, E., Fogarty, M., and Munch, S. (2012). Detecting causality in complex ecosystems. *Science*, 338, 496-500.